RESKILLING LAB
DT-01 TRAINING & SELF-LEARNING PACK

Artificial Intelligence & Generative AI

From Digital User to AI-Enabled Professional

Beginner → Intermediate | Recommended learning time: 40–60 hours
Web-ready modular training content | Version 1.0 | September 2026

How to Use This Training Pack

This training pack is designed for self-learning, blended learning and trainer-supported delivery. It combines concise theory, guided practice, workplace tasks, SME-oriented case studies, assessment and curated external learning resources. Learners should complete the Initial Skills Check first, work through Units 1–8 in sequence, complete the practical tasks, and finish with the capstone project and final assessment.

The material is vendor-neutral. Named platforms are used as optional examples and external learning resources. Where an organisation has an internal AI policy, approved-tool list or data-classification scheme, those rules take precedence over generic examples in this pack.

1. Pathway Overview

Item	DT-01 Specification
Pathway	Artificial Intelligence & Generative AI
Target level	Beginner → Intermediate
Recommended workload	40–60 hours
Target groups	Employees; SME staff; VET learners; jobseekers; career changers; trainers; administrative and business professionals
Entry requirements	Basic digital literacy; ability to use a web browser, office applications and online services
Possible target roles	AI Support Specialist; AI-enabled Business Professional; Junior AI Practitioner; AI-enabled Administrative/Marketing/Training roles
Learning mode	Self-learning, blended learning or trainer-supported
Main outputs	Prompt portfolio; AI task redesign; risk checklist; responsible-use plan; capstone AI-enabled workflow

1.1 Learning Outcomes

- Explain core concepts of artificial intelligence, machine learning, generative AI, large language models and AI agents in accessible workplace language.
- Distinguish suitable and unsuitable uses of generative AI in common SME and professional tasks.
- Create structured prompts using context, task, constraints, examples and output-format requirements.
- Critically verify AI-generated text, summaries, analyses and recommendations before business use.
- Apply data-minimisation, privacy, confidentiality, security and intellectual-property precautions when using AI tools.
- Recognise common risks including hallucination, bias, over-reliance, prompt injection, data leakage and misleading synthetic content.
- Use generative AI to support productivity tasks such as drafting, summarising, ideation, classification, planning and knowledge work.
- Design a simple human-in-the-loop AI-enabled workflow for an SME or workplace context.
- Document responsible AI-use rules appropriate to a small organisation.
- Identify further learning and credential options for continued AI upskilling.

1.2 European Skills and Policy Context

The pathway is aligned conceptually with European digital competence and AI-literacy priorities. DigComp 2.2 includes knowledge, skills and attitudes for confident, critical and safe engagement with digital technologies, including AI-driven systems. The EU AI Act also addresses AI literacy for staff and other persons dealing with AI systems on behalf of providers and deployers. This training pack is educational guidance, not legal advice.

DigComp 2.2 – Digital Competence Framework for Citizens — European Commission Joint Research Centre.
https://publications.jrc.ec.europa.eu/repository/handle/JRC128415

European Commission – AI talent, skills and literacy — Current EU information on AI literacy.
https://digital-strategy.ec.europa.eu/en/policies/ai-talent-skills-and-literacy

EU AI Act – Regulation (EU) 2024/1689 — Official legal text.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj

2. Initial Skills Check

Rate each statement before starting the pathway. Use: 0 = No, 1 = Partly, 2 = Yes. Repeat the same assessment after completing the pathway.

#	Statement	0 No	1 Partly	2 Yes
1	I can explain the difference between AI, machine learning and generative AI.	☐	☐	☐
2	I understand, at a basic level, what a large language model does.	☐	☐	☐
3	I can identify tasks where generative AI can add value.	☐	☐	☐
4	I can write a prompt with clear context, task and output requirements.	☐	☐	☐
5	I can improve a weak prompt through iteration.	☐	☐	☐
6	I verify important AI-generated information before using it.	☐	☐	☐
7	I know what information should not be entered into an unapproved public AI tool.	☐	☐	☐
8	I can recognise possible bias, hallucination or misleading output.	☐	☐	☐
9	I understand why human oversight remains necessary.	☐	☐	☐
10	I can describe basic responsible-AI principles.	☐	☐	☐
11	I can use AI to support at least three workplace tasks.	☐	☐	☐
12	I can redesign a simple workflow to include AI without removing human accountability.	☐	☐	☐

Indicative interpretation: 0–8 = Foundation needed; 9–16 = Developing; 17–20 = Functional; 21–24 = Ready for intermediate practice. This score is formative and is not a formal qualification.

Unit 1 – AI Fundamentals and the Changing Workplace

Theory

Artificial intelligence is a broad field concerned with computer systems performing tasks associated with perception, prediction, language, pattern recognition, decision support and content generation. Machine learning is a family of methods through which models learn patterns from data rather than relying only on explicitly programmed rules. Generative AI focuses on producing new outputs such as text, images, audio, code or structured information.

For SMEs, the practical question is not whether a tool appears intelligent, but whether it can improve a defined task with acceptable cost, quality and risk. High-value starting points are often repetitive knowledge tasks: first drafts, document summarisation, classification, FAQ preparation, idea generation and extraction of structured information. AI should not be treated as an autonomous authority for legal, financial, safety-critical, medical or employment decisions.

Key Concepts

- What AI is and is not
- AI, machine learning, deep learning and generative AI
- Common AI workloads
- Where AI creates value in SMEs
- Limits of AI and the role of human judgement

Practical Activities

1. Create a two-column inventory of five tasks in your current or target job: tasks where AI could assist, and tasks where AI should not make the final decision.
2. Select one repetitive information task and describe how AI could reduce time without removing human review.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 2 – Generative AI and Large Language Models

Theory

Large language models generate text by predicting likely sequences based on patterns learned during training and, in some systems, additional retrieved or supplied context. Their fluent language can create an illusion of certainty. A well-written answer can still be wrong, incomplete, outdated or unsupported. Learners therefore need to separate fluency from reliability.

An AI agent extends the basic prompt-response pattern by combining a model with goals, tools, memory or multi-step actions. Agents can increase productivity but also increase operational risk because they may take actions rather than merely suggest text. Permissions, monitoring and approval points therefore matter.

Key Concepts

- Generative models and content creation

- Large language models (LLMs) in plain language
- Tokens, context and probabilistic generation
- Why hallucinations occur
- Grounding, retrieval and source-based answers
- Introduction to AI agents

Practical Activities

3. Ask an approved AI tool to explain the same topic for three audiences: a child, an SME manager and a technical specialist. Compare vocabulary, detail and assumptions.
4. Give the tool a short source text and ask it to answer only from that text. Mark any unsupported statements.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 3 – Prompting for Reliable Results

Theory

Prompt engineering is the practical discipline of expressing a task so that an AI system has enough context and constraints to produce a useful result. A strong prompt usually states the context, objective, relevant input, constraints, desired output format and quality criteria. Examples can clarify expectations.

Prompting is iterative. The first output should be treated as a draft. Improve results by identifying what is missing, correcting assumptions, narrowing scope, supplying evidence and defining acceptance criteria. Complex tasks should often be decomposed into smaller stages.

Key Concepts

- Prompt anatomy: role/context/task/constraints/output
- Specificity and examples
- Prompt iteration
- Decomposition of complex tasks
- Structured outputs
- When not to prompt with sensitive data

Practical Activities

5. Rewrite three vague prompts into structured prompts.
6. Create a prompt that produces a table with fixed columns and explicit acceptance criteria.
7. Build a reusable prompt template for one recurring workplace task.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 4 – AI for Workplace Productivity

Theory

Generative AI can support knowledge work by reducing the effort required to produce a first draft or transform existing information. Productivity gains are strongest when the user understands the task well enough to evaluate the result. The preferred pattern is AI-assisted work, not unreviewed AI substitution.

Good candidates include summarising non-sensitive documents, drafting agendas, creating question lists, converting notes into structured tables, brainstorming alternatives and rewriting content for different audiences. Organisational policy and data sensitivity must be checked before any content is uploaded.

Key Concepts

- Drafting and rewriting
- Summarisation
- Meeting and action-note support
- Research assistance and question generation
- Classification and extraction
- Planning and brainstorming
- Translation and localisation with human review

Practical Activities

8. Turn a fictional 500-word meeting note into decisions, actions, owners and deadlines.
9. Create three versions of a customer-facing message: formal, plain-language and concise.
10. Use AI to create a first draft of a weekly work plan, then manually correct priorities and assumptions.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 5 – Responsible AI, Ethics and Human Oversight

Theory

Responsible AI means designing and using AI in ways that preserve human accountability and address foreseeable harms. Common principles include fairness, transparency, robustness, privacy, security and accountability. Human oversight is not merely a final approval click; the reviewer must have enough information, competence and authority to challenge the output.

Bias can enter through training data, prompt framing, incomplete context or the user's own assumptions. Responsible use therefore includes testing outputs across relevant groups and avoiding automated conclusions where consequences are significant.

Key Concepts

- Fairness and bias
- Transparency
- Accountability
- Human oversight
- Accessibility and inclusion

- Intellectual property and attribution
- Appropriate disclosure of AI assistance

Practical Activities

11. Review an AI-generated recruitment or customer-service text for potentially unfair assumptions.
12. Create a five-rule Responsible AI card for a small company.
13. Identify which decisions in a sample workflow must remain with a human.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 6 – Privacy, Security and Safe AI Use

Theory

Generative AI introduces familiar information-security risks in a new interface. Users may disclose personal data, trade secrets, credentials, customer information or unpublished documents simply by pasting them into a prompt. The safest default is data minimisation: provide only what is necessary and use organisation-approved tools and configurations.

Prompt injection is an instruction-manipulation risk in which untrusted content attempts to make an AI system ignore intended rules or reveal information. Users should treat external content as data to analyse, not automatically as instructions to follow.

Key Concepts

- Personal data and confidential information
- Data minimisation
- Approved versus shadow AI
- Account security and MFA
- Prompt injection and malicious content
- Sensitive files and third-party tools
- Incident reporting

Practical Activities

14. Classify ten example data items as safe, restricted or prohibited for an unapproved public AI tool.
15. Inspect a fictional prompt-injection example and explain why external instructions should not automatically override organisational rules.
16. Draft a short incident report for accidental disclosure of confidential information to an AI service.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 7 – Evaluating AI Outputs and Managing Risk

Theory

AI output evaluation should be systematic. Check factual accuracy, source support, relevance to the task, completeness, internal consistency, recency, bias, safety and usability. Important numerical claims should be recalculated or verified from authoritative sources.

A practical verification rule is: the higher the consequence of error, the stronger the evidence and human review required. AI can help locate questions and possibilities; accountability for business use remains with people and organisations.

Key Concepts

- Accuracy and hallucination
- Source verification
- Bias checks
- Freshness and completeness
- Quantitative claims
- High-impact decisions
- Verification workflow and evidence log

Practical Activities

17. Fact-check five claims in a generated answer using authoritative sources.
18. Use a simple evaluation rubric: accuracy, relevance, completeness, safety and usability.
19. Compare two AI outputs for the same business task and justify which elements you would retain, revise or reject.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

Unit 8 – Building an AI-Enabled SME Workflow

Theory

AI adoption should begin with a process, not a tool. Map the existing workflow, identify bottlenecks, define the desired outcome and then decide whether AI is suitable. Build explicit review and escalation points. Avoid automating a poorly understood process.

A simple SME pilot should have a narrow scope, an owner, approved data inputs, a review method and measurable indicators. After the pilot, compare quality, time, cost and risk with the previous process before expanding use.

Key Concepts

- Task selection
- Process mapping
- Human-in-the-loop design
- Input and output controls
- Approval points

- Performance measures
- Continuous improvement and governance

Practical Activities

20. Map a current workflow using: Trigger → Input → AI-supported task → Human review → Decision → Output → Record.
21. Define three KPIs such as time saved, error rate, rework rate or user satisfaction.
22. Prepare a one-page AI Use Procedure for the selected workflow.

Knowledge Check

Before moving on, explain the key concepts in your own words and record one workplace example, one benefit and one risk associated with this unit.

3. Prompt Engineering Toolkit

Use the following reusable structure for business prompts. Do not include confidential or personal information unless your organisation has explicitly approved the tool and use case.

Element	Question to answer	Example
Context	What situation should the AI understand?	We are a 20-person manufacturing SME preparing a staff awareness session.
Task	What exactly should it do?	Draft a 30-minute session outline.
Input	What information should it use?	Use the five bullet points below.
Constraints	What must it avoid or respect?	Use plain English; do not invent legal requirements.
Output	What format is required?	Return a table with time, topic, activity and facilitator note.
Quality check	How should the result be tested?	Flag any statement that needs external verification.

3.1 Reusable Prompt Templates

Use case	Template
Summarisation	Summarise the following material for [audience]. Preserve the original meaning. Separate: key points, decisions, actions and unresolved questions. Do not add facts not present in the source.
Comparison	Compare [A] and [B] using only these criteria: [criteria]. Present similarities, differences, risks and unanswered questions. Do not select a winner unless explicitly requested.
Drafting	Draft a [document type] for [audience] with the objective [objective]. Tone: [tone]. Maximum length: [length]. Use only the supplied facts. Mark missing information as [INFORMATION NEEDED].
Extraction	Extract the following fields from the supplied text: [fields]. Return a table. If a field is absent, write 'Not stated'.
Critical review	Review the draft for factual claims, unsupported assumptions, ambiguity, privacy/security concerns and possible bias. Return a correction table; do not rewrite until the issues are listed.

4. SME Case Studies

Case 1 – Customer Support Drafting

A 12-person online retailer receives repetitive customer questions. Staff want AI to draft responses. Design a safe workflow in which customer personal data is minimised, approved templates are used, and a human checks refunds, complaints and legally sensitive responses.

- Which data should be removed before prompting?
- Which responses can be low-risk drafts?
- Which cases require escalation?
- How will quality be monitored?

Case 2 – Marketing Content

A small tourism company wants to generate multilingual social-media posts and images. The team is under pressure to publish quickly.

- How will factual claims be verified?
- How will copyright, brand and disclosure issues be handled?
- What human approval is required before publication?
- How can the team test whether translated content is culturally appropriate?

Case 3 – HR Screening Risk

An SME considers using generative AI to summarise CVs and suggest interview candidates.

- What fairness and privacy risks arise?
- What information should not be inferred?
- Why should the tool not make the final hiring decision?
- What audit trail and human review would be needed?

5. Practical Lab Series

Lab	Topic	Task	Evidence
Lab 1	Prompt Improvement	Turn a vague prompt into a structured prompt and document three iterations.	Prompt log + final prompt
Lab 2	Source-Grounded Summarisation	Summarise a supplied non-sensitive document and identify unsupported additions.	Annotated summary
Lab 3	Structured Extraction	Extract actions, dates and owners from fictional meeting notes.	Action table
Lab 4	Content Transformation	Rewrite one text for customer, manager and learner audiences.	Three versions + comparison
Lab 5	Risk Review	Evaluate a generated answer for accuracy, bias, privacy and security.	Risk checklist
Lab 6	Workflow Redesign	Add AI assistance to a current SME process while preserving human review.	Process map

Lab 7	Responsible AI Rules	Draft a one-page internal AI-use guide.	AI-use guide
Lab 8	Capstone Prototype	Demonstrate an AI-enabled workflow and evaluate it.	Portfolio + presentation

6. Capstone Project – AI-Enabled SME Workflow

Design and document a small, realistic AI-enabled workflow for an SME, VET provider or workplace. The project must demonstrate productivity value and responsible use.

Required Deliverables

23. Problem statement and current workflow.
24. Target users and business objective.
25. Selected AI-assisted task and justification.
26. Three tested prompts with iteration notes.
27. Data classification and privacy/security precautions.
28. Human review and escalation points.
29. Example AI output with documented verification.
30. Responsible-use rules for staff.
31. Three performance indicators.
32. One-page reflection: benefits, limitations, risks and next steps.

Assessment Rubric

Criterion	Weight	Evidence
Problem and workflow definition	15%	Clear need, scope and users
Prompt/task design	20%	Structured prompts and iteration
Practical usefulness	20%	Relevant, usable workplace output
Verification and human oversight	20%	Evidence checks and approval points
Privacy, security and responsible AI	20%	Appropriate safeguards
Reflection and communication	5%	Clear limitations and next steps

7. Final Knowledge Assessment

1. Which statement best describes generative AI?

A. It only stores files
B. It creates new outputs based on learned patterns and context
C. It guarantees factual answers
D. It replaces all human judgement

2. Why can an LLM produce a fluent but incorrect answer?

A. Fluency is not the same as verified factual accuracy
B. It always searches official databases
C. It refuses uncertain tasks
D. It only repeats exact source text

3. Which prompt element most directly specifies the desired structure of the answer?

A. Context
B. Output format
C. Role
D. Background

4. What is the safest default for sensitive business data?

A. Paste everything for context
B. Use any free tool
C. Minimise data and use approved tools/policies
D. Remove only the company logo

5. What is hallucination?

A. A verified source
B. A plausible but unsupported or incorrect generated statement
C. A firewall alert
D. A password attack

6. What does human-in-the-loop mean?

A. A person meaningfully reviews or controls relevant AI-supported steps
B. AI works without oversight
C. Humans only buy the software
D. AI decisions are always final

7. Which is a good verification practice?

A. Trust polished wording
B. Verify important claims with authoritative evidence
C. Ask the same model whether it is correct and stop there
D. Ignore dates

8. What is prompt injection?

A. A method of compressing files
B. Untrusted instructions attempting to manipulate AI behaviour
C. A type of spreadsheet
D. A password manager

9. Which is the strongest first step for SME AI adoption?

A. Buy the most expensive model
B. Automate every process
C. Define a specific process problem and acceptable risk
D. Remove human review

10. Why should organisations document AI-use practices?

A. To support governance, consistency, learning and accountability
B. To make prompts longer
C. To eliminate all risk
D. To avoid staff training

Answer Key

1-B, 2-A, 3-B, 4-C, 5-B, 6-A, 7-B, 8-B, 9-C, 10-A.

8. Post-Training Skills Check

Repeat the Initial Skills Check. Compare the pre- and post-training scores. More important than the total score is the learner's evidence: prompt portfolio, verification records, practical labs and capstone workflow.

Reflection Questions

- Which AI-supported task can you now perform more effectively?
- Which AI risk are you most likely to encounter in your work?
- What information will you never enter into an unapproved AI tool?
- What verification steps will you apply before using AI-generated content?
- Where must human decision-making remain central in your workplace?
- What is your next learning objective?

9. Recommended Tools and Practice Environment

Tool availability, pricing, privacy terms and organisational approval can change. Learners should use tools permitted by their organisation and review current terms before uploading data.

Purpose	Examples / approach	Practice rule
Conversational generative AI	Use an organisation-approved general-purpose assistant	Use fictional or non-sensitive data for training
Document summarisation	Approved AI assistant or local training text	Check every key claim against the source
Image generation	Approved image-generation service	Check licensing, branding and disclosure needs
Office productivity	AI features integrated into approved office suite	Do not assume enterprise privacy controls apply to personal accounts
Research assistance	AI plus authoritative web sources	Use AI to support discovery; verify with primary sources
Workflow prototyping	No-code/automation or manual simulation	Keep a human approval step during training

10. Curated Learning Resources

The following resources were selected because they provide authoritative or structured learning relevant to this pathway. Availability and certification conditions should be checked on the provider's current page.

Resource	URL	Recommended use
European Commission JRC	https://publications.jrc.ec.europa.eu/repository/handle/JRC128415	European digital competence framework;

– DigComp 2.2		includes AI-related examples.
European Commission – AI Literacy	https://digital-strategy.ec.europa.eu/en/policies/ai-talent-skills-and-literacy	EU AI-literacy policy information and implementation context.
European Commission – AI Literacy Q&A	https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers	Practical explanation of Article 4 and organisational AI-literacy actions.
EU AI Office – Repository of AI Literacy Practices	https://digital-strategy.ec.europa.eu/en/policies/repository-ai-literacy-practices	Examples of organisational AI-literacy initiatives.
EUR-Lex – EU AI Act	https://eur-lex.europa.eu/eli/reg/2024/1689/oj	Official legal framework; use for authoritative legal reference.
OECD AI Principles	https://oecd.ai/en/ai-principles	International principles for trustworthy AI.
Microsoft Learn – Introduction to generative AI and agents	https://learn.microsoft.com/en-us/training/modules/fundamentals-generative-ai/	Beginner module on LLMs, prompts and agents.
Microsoft Learn – Build foundational generative AI skills	https://learn.microsoft.com/en-us/training/paths/explore-microsoft-365-copilot-business-chat/	Beginner learning path covering responsible AI, prompting and productivity.
Microsoft – Generative AI for Beginners	https://learn.microsoft.com/en-us/shows/generative-ai-for-beginners/	Structured beginner series with multiple lessons.
IBM SkillsBuild	https://skillsbuild.org/	Free learning platform with AI and generative-AI learning experiences; account conditions may apply.
Elements of AI	https://course.elementsofai.com/	Free online course covering

		AI concepts, problem solving, machine learning, neural networks and implications.
Google – AI for All / Herkes için Yapay Zeka	https://grow.google/intl/tr/ai-for-all/	AI learning resources, including Turkish-language introductory content.

11. Certification and Microcredential Path

ReskillingLab should distinguish between (a) completion of this internal pathway, (b) third-party course completion records/badges, and (c) formal industry certifications. A course-completion badge should not be presented as equivalent to a regulated qualification or professional certification.

- ReskillingLab pathway completion: issue only after the required labs, final assessment and capstone evidence are completed.
- Microsoft Learn: selected modules provide profile-based completion/achievement records; current conditions vary by module.
- IBM SkillsBuild: selected learning experiences may provide digital credentials; check the current credential description before enrolment.
- Elements of AI: completion/credit arrangements can depend on course version and learner eligibility.
- Google learning programmes: certificate conditions vary by course and region; verify the current offer before advertising a credential.

12. Career Transition Guide

Starting profile	**Possible next role direction**	**Evidence to build**
Administrative / office staff	AI-enabled administrative professional	Prompt portfolio; meeting/action extraction; workflow redesign
Marketing / communications	AI-enabled content or marketing assistant	Content workflow; verification checklist; brand-safe prompts
Trainer / VET professional	AI-enabled trainer / learning-content developer	Lesson design; assessment prompts; responsible-use guide
SME owner / manager	AI adoption coordinator / AI-enabled manager	Use-case inventory; risk controls; KPI-based pilot
IT support / digital staff	Junior AI support / AI operations pathway	AI fundamentals; security practices; workflow and agent concepts
Jobseeker / career changer	Entry-level AI-enabled business role	Capstone portfolio; documented practical tasks; external learning

		records

Suggested next learning directions: advanced prompt engineering; data analytics; automation; AI agents; retrieval-augmented generation; AI governance; cybersecurity for AI; or role-specific AI applications.

13. Glossary

Term	Definition
AI	Artificial Intelligence: computer-based systems performing tasks associated with capabilities such as prediction, perception, language or decision support.
Machine Learning	Methods that learn patterns from data to make predictions or generate outputs.
Generative AI	AI capable of producing new content such as text, images, audio or code.
LLM	Large Language Model; a model trained on large amounts of text and related data to process and generate language.
Prompt	Input or instruction provided to a generative AI system.
Token	A unit of text processed by many language models.
Context Window	The amount of information a model can consider within an interaction.
Hallucination	A generated statement that appears plausible but is incorrect or unsupported.
Grounding	Connecting model output to supplied or retrieved evidence.
RAG	Retrieval-Augmented Generation; a pattern combining information retrieval with generation.
AI Agent	An AI-enabled system that can pursue goals and may use tools or execute multi-step actions.
Human-in-the-loop	A design in which people meaningfully review, approve or control relevant AI-supported steps.
Bias	Systematic skew that can produce unfair or distorted outcomes.
Prompt Injection	Malicious or untrusted instructions intended to manipulate an AI system's behaviour.
Data Minimisation	Using only the data necessary for a defined purpose.
Shadow AI	Use of AI tools outside an organisation's approved governance or IT process.
MFA	Multi-Factor Authentication; account protection requiring more than one authentication factor.
AI Literacy	Knowledge, skills and understanding needed to engage with AI appropriately in a given context.

14. Trainer / Facilitator Notes

- Use fictional, anonymised or explicitly approved datasets in demonstrations.
- Ask learners to show their verification process, not only the final AI output.
- Reward identification of limitations and risks; do not reward confident but unsupported AI-generated text.
- Adapt examples to the learner's sector while keeping privacy and safety constraints.

- Where learners use different AI tools, assess the process and evidence rather than brand-specific features.
- Keep legal and policy examples current; distinguish educational guidance from legal advice.
- Encourage learners to maintain a prompt-and-evidence portfolio throughout the pathway.

15. Suggested 40–60 Hour Delivery Plan

Component	Indicative hours
Initial skills check and orientation	2
Unit 1 – AI Fundamentals	5
Unit 2 – Generative AI and LLMs	5
Unit 3 – Prompting	7
Unit 4 – Workplace Productivity	7
Unit 5 – Responsible AI	5
Unit 6 – Privacy and Security	5
Unit 7 – Output Evaluation and Risk	5
Unit 8 – SME Workflow Design	5
Capstone, assessment and reflection	4–14
Total	50–60 (adaptable to a 40-hour accelerated route)

16. References

European Commission, Joint Research Centre. (2022). DigComp 2.2: The Digital Competence Framework for Citizens. Publications Office of the European Union. https://doi.org/10.2760/115376

European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). EUR-Lex.

European Commission. (2026). AI talent, skills and literacy. Shaping Europe's digital future.

European Commission. (2026). AI Literacy – Questions & Answers. Shaping Europe's digital future.

OECD. (2024). OECD AI Principles (updated 2024). OECD.AI.

Microsoft Learn. Introduction to generative AI and agents.

Microsoft Learn. Build foundational generative AI skills.

University of Helsinki / MinnaLearn. Elements of AI.

IBM SkillsBuild. Artificial intelligence learning resources.

Google. AI for All / Herkes için Yapay Zeka.

17. Website Publishing Structure

For ReskillingLab, publish this document as modular web content rather than one long page. Recommended navigation:

- Overview
- Initial Skills Check
- Units 1–8
- Prompt Toolkit
- Case Studies
- Practical Labs
- Capstone Project
- Final Assessment
- Learning Resources
- Credentials
- Career Transition
- Glossary

Each unit should have four visible tabs or blocks: Learn | Practice | Check Your Understanding | Resources. Learner progress can later be connected to login/profile functionality if the platform supports it.

18. Expanded Theoretical Learner Handbook

This section expands the theoretical foundation of DT-01. It is intended to be read alongside Units 1–8 and can be published on ReskillingLab as extended learning content. The emphasis is on conceptual understanding, workplace relevance, limitations and responsible application.

18.1 Artificial Intelligence: Concepts, Scope and Evolution

Artificial intelligence is an umbrella term for computational systems that perform functions commonly associated with human cognitive activity, such as recognising patterns, interpreting language, predicting outcomes, generating content and supporting decisions. AI is not a single technology. It includes rule-based systems, statistical learning, machine learning, deep learning, computer vision, natural-language processing, recommender systems, generative models and increasingly agentic systems that can plan and execute multi-step tasks.

A useful distinction is between traditional software and learning-based AI. Traditional software follows rules explicitly written by developers: given a defined input, it applies programmed logic and produces an output. Machine-learning systems instead learn statistical relationships from examples. Their behaviour therefore depends not only on source code but also on training data, model architecture, optimisation procedures and the context in which the model is deployed. This is one reason AI systems require different forms of testing, monitoring and governance from conventional software.

Machine learning can be broadly divided into supervised, unsupervised and reinforcement learning. Supervised learning uses labelled examples to learn a relationship between inputs and expected outputs, for example classifying messages as spam or legitimate. Unsupervised learning searches for patterns or structure without predefined labels, for example grouping customers with similar purchasing behaviour. Reinforcement learning involves an agent learning through interaction with an environment and feedback related to its actions. Modern AI products may combine several approaches.

Deep learning is a branch of machine learning based on multi-layer neural networks. It has been especially influential in image recognition, speech processing and language technologies. Large-scale neural models can learn complex representations from vast datasets, but their performance comes with trade-offs: high computational demand, limited interpretability, dependence on training data quality and the possibility of unexpected behaviour outside familiar conditions.

For workplace users, the most important lesson is that AI should be understood as a capability embedded in a process rather than as an independent source of truth. A system can be highly effective at one task and unreliable at another. Appropriate use therefore starts by defining the task, expected output, acceptable error rate, data sensitivity and consequences of failure.

18.2 From Predictive AI to Generative AI

Predictive AI estimates a category, probability or numerical outcome from available information. Examples include fraud detection, demand forecasting, predictive maintenance and customer-churn prediction. Generative AI has a different emphasis: it produces new content or representations. Depending on the model, that content may include text, images, audio, video, software code, synthetic data or combinations of several media types.

Generative AI does not normally retrieve a pre-written answer from a database. A language model generates a response dynamically from patterns represented in its parameters and from the context supplied during the interaction. This distinction explains both its flexibility and its risk. The same system can draft a report, translate a paragraph, classify comments and generate code, but it may also create unsupported details because generation is not identical to factual retrieval.

Foundation models are large models trained broadly enough to support many downstream tasks. They can be adapted through prompting, fine-tuning, retrieval systems or tool connections. This has lowered the barrier to AI adoption: an SME may gain useful capabilities without training a model from scratch. At the same time, organisations become dependent on decisions made by external providers concerning model design, data processing, security, availability and product terms.

Multimodal AI can process or generate more than one type of information, such as text and images, or text, audio and video. In a business context this enables tasks such as analysing a photographed document, describing a diagram, transcribing a meeting, generating marketing visuals or combining written instructions with screenshots. Multimodality increases utility but also expands the range of privacy, copyright, authenticity and security questions that must be considered.

18.3 How Large Language Models Work

A large language model processes text as tokens. A token is not necessarily a whole word; it can be a word, part of a word, punctuation or another textual unit. During training, the model learns statistical relationships among tokens across very large collections of text. During generation, it estimates which token is likely to follow the preceding context and repeats this process to construct a response.

This mechanism helps explain why an LLM can produce coherent language without possessing human understanding in the ordinary sense. It can model linguistic and conceptual relationships with remarkable sophistication, yet it does not automatically know whether a statement is true, current or appropriate for a specific organisation. It also does not have direct access to every source that may have influenced training.

The context window is the amount of information the model can consider during a particular interaction. Instructions, conversation history, uploaded content and retrieved material may all consume context. Larger context windows allow more material to be considered, but more context does not automatically produce a better answer. Irrelevant, contradictory or poorly structured material can reduce performance.

Temperature and related sampling settings influence how deterministic or varied generation may be, although many consumer interfaces hide these controls. For tasks requiring consistency, such as extraction or classification, lower variability is usually desirable. For ideation and creative exploration, greater variation can be useful. Users should nevertheless judge outputs by task-specific quality criteria rather than by a single technical setting.

An important limitation is hallucination: the generation of content that is false, fabricated or unsupported while sounding plausible. Hallucination cannot be solved merely by asking the model to 'be accurate'. It is managed through better task design, authoritative source material, retrieval, explicit uncertainty handling, verification and human review.

18.4 Retrieval, Grounding and Knowledge-Augmented AI

Grounding means connecting an AI response to defined evidence rather than relying only on the model's internal statistical representations. A simple form of grounding occurs when a user supplies a document and instructs the model to answer only from that document. More advanced systems use retrieval-augmented generation (RAG), in which relevant passages are retrieved from an external knowledge base and supplied to the model when a question is asked.

RAG is attractive to organisations because internal policies, manuals, product information or project documents can remain separate from the base model while still informing answers. A typical RAG workflow includes document ingestion, text segmentation, representation or indexing, retrieval of relevant passages and generation based on those passages. The quality of each stage affects the final answer.

Grounding reduces some forms of hallucination but does not eliminate error. The system may retrieve the wrong passage, omit an important source, misunderstand the question, misinterpret evidence or combine correct fragments incorrectly. Users therefore need citations, traceability and a method for checking the underlying source where consequences are significant.

For SMEs, a modest and well-curated knowledge base is often more valuable than a large unstructured collection of files. Document ownership, version control, retention, access permissions and removal of obsolete information should be addressed before AI is used as an interface to organisational knowledge.

18.5 Prompt Engineering as Task Specification

Prompt engineering is best understood as task specification rather than as a collection of magic phrases. A useful prompt communicates what the user wants, why it is needed, what information is authoritative, what constraints apply and how the output should be structured. The clearer the task specification, the easier it is to evaluate whether the response meets the requirement.

Context explains the situation. The task states the action. Constraints define boundaries such as length, audience, tone, prohibited assumptions or source limitations. Output specifications define structure, for example a table with fixed columns. Quality criteria tell the model and the human reviewer what a satisfactory result should contain. When relevant, examples can demonstrate the expected pattern.

Prompt decomposition is valuable for complex work. Instead of asking for a complete strategy in one request, the user can separate the process into stages: identify facts, identify missing information, generate alternatives, evaluate alternatives against criteria, draft the output and finally perform a critical review. This makes errors easier to detect and reduces the risk that an early assumption silently propagates through the whole output.

Few-shot prompting provides examples of desired inputs and outputs. It can improve consistency for classification, formatting and transformation tasks. However, examples may also introduce unintended bias if they are not representative. Prompt templates should therefore be tested with varied cases rather than accepted after a single successful output.

A mature prompting practice includes versioning and evidence. For recurring organisational tasks, record the prompt version, purpose, permitted inputs, expected outputs, verification steps and known limitations. This turns prompting from an individual trick into a manageable organisational process.

18.6 Generative AI for Knowledge Work and Productivity

Generative AI can reduce the time required for drafting, restructuring and exploring information. It is especially useful where a knowledgeable worker can quickly judge whether the output is suitable. Examples include drafting agendas, converting notes into action lists, producing alternative formulations, extracting fields from text, generating interview questions, creating first-pass summaries and adapting content for different audiences.

Productivity should not be measured only as time saved. If AI produces more text but increases checking and correction, apparent speed may not translate into real efficiency. Useful metrics include cycle time, error rate, rework, consistency, user satisfaction and the proportion of outputs accepted after human review. A pilot should compare these measures with the previous process.

Automation bias occurs when people give excessive weight to a system's recommendation because it appears objective or technologically sophisticated. Generative AI can intensify this effect because explanations are fluent and confident. Training should therefore cultivate active review: users should ask

what evidence supports the answer, what assumptions were made, what may be missing and what would happen if the answer were wrong.

AI can also change job design. Some tasks may be partially automated, while other tasks become more important: problem formulation, verification, domain judgement, stakeholder communication, ethical review and workflow design. Reskilling should therefore combine technical tool use with critical thinking and domain competence.

18.7 Responsible AI: Fairness, Transparency and Accountability

Responsible AI is the disciplined effort to obtain value from AI while respecting people, rights, organisational responsibilities and foreseeable risks. It is not achieved by a single checklist. It requires governance, competence, documentation, technical safeguards and appropriate human decision-making.

Fairness concerns whether AI-supported processes create unjustified disadvantages for individuals or groups. Bias can arise from historical data, incomplete representation, proxy variables, model design, prompt wording or the way outputs are interpreted. A fairness review should examine the full process, not merely the final model output.

Transparency has several levels. Users may need to know that AI is being used, decision-makers may need to understand the basis and limitations of the system, and technical teams may need logs or documentation sufficient to investigate incidents. The appropriate level depends on context and risk. Transparency does not mean revealing confidential model details; it means providing meaningful information to the people who need it.

Accountability remains with people and organisations. Delegating a task to an AI service does not automatically transfer responsibility for the consequences. A sound governance model identifies the process owner, permitted uses, review requirements, escalation routes and responsibility for monitoring.

Human oversight must be substantive. A person who routinely approves outputs without adequate time, expertise or access to evidence is not providing meaningful oversight. Organisations should design review points so that a human can understand the decision context, challenge the output and intervene before harm occurs.

18.8 Privacy, Confidentiality and Data Governance

AI tools can receive large amounts of information through prompts, files, images and integrations. Before entering information, users should determine its classification: public, internal, confidential, personal, special-category personal data, commercially sensitive or otherwise restricted. Organisational policy should specify which categories may be used with which services.

Data minimisation means providing only the information necessary for the task. Names can sometimes be replaced with roles or identifiers; unnecessary columns can be removed from datasets; examples can be fictionalised; and sensitive documents can be summarised locally before an approved subset is provided to an AI service. Minimisation reduces exposure even when a tool has strong security controls.

Users should distinguish between a provider's public consumer service and its enterprise or organisational offering. Data handling, retention, administrative controls and contractual commitments may differ. Assumptions about one product tier should not automatically be applied to another.

AI governance should also address retention and lifecycle. An organisation needs to know where prompts, uploaded files, generated outputs and logs are stored, who can access them, how long they remain available and how they are deleted. These questions become more important when AI is connected to email, document repositories or customer systems.

18.9 Security Threats in Generative AI Use

Generative AI introduces new attack surfaces while also interacting with familiar cybersecurity risks. Account compromise can expose conversation histories or connected data. Malicious files or websites may contain instructions designed to manipulate an AI system. Over-permissioned agents may be able to send messages, modify files or trigger external actions beyond what the user intended.

Prompt injection occurs when untrusted content contains instructions that compete with the intended system or user instructions. For example, a document being summarised might contain text telling the AI to ignore previous rules or reveal confidential information. The correct defensive principle is to separate data from authority: content being analysed should not automatically be treated as trusted instructions.

Indirect prompt injection is especially relevant when an AI system retrieves web pages, emails or documents automatically. The user may never see the malicious instruction. Defences therefore require system design: restricted permissions, isolation of untrusted content, confirmation before consequential actions, monitoring and limits on what tools an AI agent can access.

AI-generated code and scripts require normal secure-development practices. Code that appears plausible may contain vulnerabilities, insecure dependencies or incorrect assumptions. Human review, testing, dependency checks and controlled deployment remain necessary.

18.10 Evaluating AI Output Quality

Evaluation should be tied to the intended use. A brainstorming output may be acceptable if it is diverse and stimulating, while a compliance summary requires precise sourcing and careful verification. Before using AI, define what quality means for the task.

Accuracy asks whether statements are correct. Relevance asks whether the response addresses the actual task. Completeness asks whether important elements are missing. Consistency checks whether parts of the answer contradict one another. Traceability asks whether important claims can be connected to evidence. Safety considers whether the output creates unacceptable legal, ethical, privacy, security or operational risk.

Quantitative information deserves special care. AI may make arithmetic mistakes, confuse units, combine incompatible datasets or invent percentages. Recalculate important figures and check the original source. For current information, verify the date and population to which a statistic applies.

A useful evaluation process separates generation from review. First produce the draft. Then review it using a checklist or rubric. For high-consequence work, use independent authoritative evidence rather than relying on the same AI system to certify its own answer.

18.11 AI Agents, Automation and Human-in-the-Loop Design

An AI agent is generally understood as a system in which an AI model can pursue a goal through multiple steps and may use external tools. An agent might search documents, update a record, draft an email and request approval before sending it. This differs from a simple chatbot that only returns text.

Agentic capability creates a spectrum of autonomy. At one end, AI suggests an action and a person performs it. Next, AI prepares an action for approval. At higher autonomy, AI may execute actions within defined limits and report afterwards. The appropriate level depends on reversibility, data sensitivity, financial impact, safety and the organisation's ability to monitor the process.

Least privilege is a key design principle. An agent should have only the permissions needed for its task. A summarisation assistant does not need permission to delete documents. A scheduling assistant may need

calendar access but not unrestricted access to financial systems. Permissions should be reviewed as workflows change.

Human-in-the-loop design identifies where judgement and approval are required. Human-on-the-loop models allow automation to proceed while a person monitors and can intervene. Fully autonomous operation should be reserved for narrowly defined, low-risk and well-monitored tasks where the consequences of error are acceptable.

18.12 AI Adoption in SMEs: From Experiment to Governance

SMEs often begin AI adoption informally, with individual employees experimenting with public tools. This can generate useful innovation but also creates shadow AI: unapproved use that bypasses normal procurement, security and data-governance processes. A practical governance approach should make safe experimentation easier rather than simply prohibiting all use.

A useful adoption sequence is: identify a business problem; classify the data and risk; select an approved tool; run a limited pilot; define human review; measure results; document lessons; and decide whether to scale. This approach prevents technology enthusiasm from replacing business analysis.

Use-case prioritisation can consider expected value, implementation effort, data readiness and risk. A low-risk, high-frequency drafting task may be a better first project than an ambitious autonomous decision system. Early wins can build organisational competence while governance matures.

An SME AI policy can remain concise. It should state approved tools, prohibited data, permitted use cases, review responsibilities, disclosure expectations, incident reporting and ownership of AI-generated work. The policy should be supported by examples and training so that employees can apply it in real situations.

AI literacy is role-dependent. A general employee needs safe-use skills, a manager needs governance and risk understanding, an IT professional needs integration and security competence, and a person overseeing high-impact uses may need deeper legal and technical knowledge. Training should therefore be proportional to responsibility.

18.13 Generative AI, Intellectual Property and Content Authenticity

Generative AI complicates familiar intellectual-property questions because outputs may be influenced by patterns learned from large training corpora, while users may also provide copyrighted or confidential source material. Organisations should avoid assuming that generated content is automatically free of third-party rights or automatically eligible for the same protection as human-created work in every jurisdiction.

A practical content workflow records the origin of source material, avoids uploading material without appropriate permission, checks generated content for obvious copying or trademark issues, and applies human editorial contribution. Where attribution is required, AI should not be relied upon to invent citations; references should be verified against the original source.

Synthetic media also creates authenticity concerns. Generated images, audio and video can be mistaken for documentary evidence. Organisations should consider labelling, provenance and approval rules, particularly for public communication. The more realistic the synthetic content, the greater the need to avoid misleading audiences.

18.14 Future Skills for an AI-Enabled Workforce

AI literacy is becoming part of general digital literacy, but it does not replace foundational skills. Reading comprehension, numeracy, domain knowledge, information literacy, cybersecurity awareness and communication remain essential because they enable people to judge AI outputs. Weak domain knowledge makes it harder to recognise a plausible error.

The AI-enabled worker needs three complementary capabilities: task competence, AI interaction competence and evaluation competence. Task competence means understanding the work itself. AI interaction competence means knowing how to frame tasks, provide context and use appropriate tools. Evaluation competence means checking evidence, risk and suitability before acting.

Career resilience also depends on learning how to redesign work. Rather than asking only 'Can AI do this task?', workers should ask 'Which parts can AI assist, which parts require human judgement, and what new responsibilities emerge?' This perspective supports reskilling rather than simple substitution.

A strong personal development plan therefore combines hands-on AI practice with a portfolio of evidence. Learners should retain examples of prompt improvement, verified outputs, workflow maps, risk assessments and reflections. These artefacts demonstrate applied competence more effectively than tool familiarity alone.

18.15 Theoretical Review Questions

1. Why is AI better understood as a family of technologies rather than a single tool?

2. How does machine learning differ from conventional rule-based software?

3. Why can a language model produce fluent text without guaranteeing factual accuracy?

4. What is the relationship between grounding, retrieval and verification?

5. Why should prompt engineering be treated as task specification?

6. How can productivity gains from generative AI be measured beyond simple time savings?

7. What makes human oversight meaningful rather than symbolic?

8. How does data minimisation reduce AI-related privacy and security risk?

9. What is indirect prompt injection and why is it relevant to AI agents?

10. How should evaluation criteria change between brainstorming and high-consequence tasks?

11. Why does greater agent autonomy require stronger permission controls?

12. What is shadow AI and how can an SME respond constructively?

13. What intellectual-property and authenticity questions should be considered before publishing AI-generated content?

14. Which capabilities make a worker more resilient in an AI-enabled labour market?